arXiv:1606.01481v1 [cs.CV] 5 Jun 2016

Better Image Segmentation by Exploiting Dense Semantic Predictions

Qiyang Zhao, Lewis D Griffin

Beihang University & UCL

Abstract. It is well accepted that image segmentation can benefit from utilizing multilevel cues. The paper focuses on utilizing the FCNN-based dense semantic predictions in the bottom-up image segmentation, arguing to take semantic cues into account from the very beginning. By this we can avoid merging regions of similar appearance but distinct semantic categories as possible. The semantic inefficiency problem is handled. We also propose a straightforward way to use the contour cues to suppress the noise in multilevel cues, thus to improve the segmentation robustness. The evaluation on the BSDS500 shows that we obtain the competitive region and boundary performance. Furthermore, since all individual regions can be assigned with appropriate semantic labels during the computation, we are capable of extracting the adjusted semantic segmentations. The experiment on Pascal VOC 2012 shows our improvement to the original semantic segmentations which derives directly from the dense predictions.

1 Introduction

People have realized the importance of utilizing multilevel cues in image segmentation for quite long a time [1][2]. The usual cues we can exploit are appearances in the low level, contours in the mid-level and semantic cues in the high level. In the widely-adopted image segmentation methods, the low and mid- level cues are exploited more frequently. It is sensible as the segmentation task is to partition images into regions of uniform perceptual features [3]. Segmentation methods developed in light of graph theory [4], nonparametric clustering [5], information theory [6], correlation clustering [7] and manifold embedding [8], exploit the appearance cues in diverse ways. Recent researches attempt to learn mid-level cues such as contours from human segmentations [9][10][11][12][13][14], then transform contour maps into segmentation hierarchies using the Ultrametric Contour Map (UCM) method [15][16][17][18]. Rather than learning from only the final human segmentations, it is suggested in [19] to learn from human behaviors conducted in the annotation processes.

However, some mistakes might be made if only appearance and contour cues are exploited, because it is hard to discriminate the regions of similar appearances but distinct semantic categories. A good example is the animals of protective colorations in the wild fields. Therefore there are a number of literatures arguing that the semantic cues should be exploited [2][20][11], optionally as the high-level features [21][22][23]. It should be noticed that although multilevel cues are involved in these approaches, the cues are utilized asynchronously rather than simultaneously. To be detailed, image pixels are grouped into superpixels according to their appearances, then the specific models are

established to exploit semantic cues [24][23].The advantage is, people can adopt different well-developed tools such as conditional random fields (CRF) in individual stages to develop segmentation algorithms easily. However as discussed in Sec. 3, postponing to exploit semantic cues can hardly avoid the similar-apperance-distinct-category mistakes though. An exception is [20] which uses a joint embedding of semantic cues and appearance cues, but there we have to solve the Normalized Cuts [25] with the huge computational load, and it is not flexible enough to accommodate the emerging dense semantic cues.

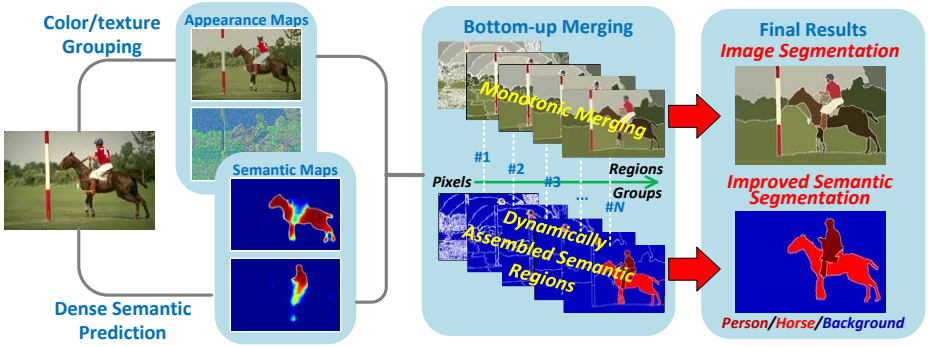

Fig. 1. The pipeline. Given an image, appearance maps and semantic maps are produced to calculate region description complexities. The bottom-up segmentation process generates a segmentation hierarchy together with a sequence of varying category-wise groups. Finally by thresholding we obtain the image segmentation, optionally with the adjusted semantic segmentation.

The semantic segmentation methods can produce semantic cues of fine granularities for arbitrary images. The precursor approaches obtain category-wise groups by refining object detection boxes [26][27][28], or classifying superxiels [16] and object proposals [29][30] into different semantic categories [31][32]. They deal with a severely limited number of semantic categories, and the segmentation precisions are usually uncompetitive. Fortunately, people have made significant progresses based on Fully-Convolutional Neural Network (FCNN) [33] in recent years. The FCNN-based semantic segmentations can be improved further [34][35][36][37] using existing tools such as fully connected conditional random fields [38]. The joint training of FCNN together with CRF will bring more benefit to the precisions [39]. The FCNN-based methods are of the state-of-the-art performance, thus more qualified as a reliable source of semantic cues. However so far as we know, few image segmentation approaches are developed to exploit the emerging FCNN-based dense semantic predictions.

In Sec. 3, we propose to utilize the multilevel cues including the dense FCNN-based semantic cues, in the framework of [19]. Semantic cues are taken into account from the very beginning of the region merging processes. Since there are usually unexplored categories beyond the capacity of even FCNN-based semantic models, we propose an

approximation method to handle the semantic inefficiency problem in Sec. 4. Besides this, the robustness of multilevel cues is suggested to play a key role in the quality of image segmentations [8][40], thus we propose a straightforward way in Sec. 5 to use contour cues to suppress the noise in multilevel cues. Finally we perform experiments to test our new methods on BSDS500 [1] and Pascal VOC12 [41] in Sec. 6.

2 Related Work

In the bottom-up merging segmentation approaches [16][17], it is beneficial to use contour cues learned from human annotations. Then rather than treating human segmentations as only static samples, what other stuff can we learn? In [19], it is noticed that although human segmentations are of the high quality, there are surprisingly many boundary mistakes, even in the renown BSDS500 dataset, see Fig. 2.a-b. Easy to see, these mistakes are made intentionally instead of by casual operations.

It reminds us there are some subtle facts which are not discovered in the human annotations. Here it should be clarified that we do not aim to teach machines reproduce boundary mistakes, but are interested in the underlying principles which bring us high quality segmentations. The above mistakes are just the by-products. It is argued in [19] that, the human annotation processes are subject to *the least effort principle* (LEP) which governs a wide range of human behaviors: a human *will strive to solve his problem in such a way as to minimize the total work that he must expend* [42]. When annotating segmentations, the effort includes understanding images in brains and figuring out the boundaries by input devices. The former is hard to model, thus only the later one is investigated in [19]. When tracing the boundaries, human subjects are inclined to choose simpler paths to save their effort, such as in Fig. 2.a-b. The tracing effort T is considered to be the product of (subjective) *unit tracing cost* λ and the tracing load L which is measured by time consumptions. Based on the observations in boundary tracing experiments supported by the annotation tool [43] (see Fig. 2.c-d), the tracing load are estimated by the mixture of experts in [19].

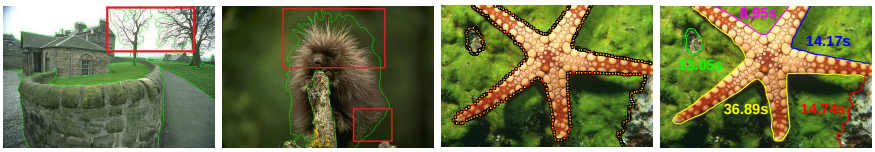

(a) (b) (c) (d)

Fig. 2. Boundary mistakes in human segmentations and the tracing experiments. In (a)(b), boundaries are shown in green and the inaccurate parts are bounded by red boxes; (c) tracing paths marked by mouse clicks of small yellow squares; (d) boundaries traced in (c) and time consumptions in seconds (in the same colors). The images and data are from [19].

Here the problem human subjects need solve is to partition images into regions of single things [1]. This target can be formulated in terms of region description complex-

ities: the lower the complexity is, the more likely only one single thing is in a region [3][44]. However it is ambiguous though for the case of moderate complexities. Generally, there is a decrease of the total description complexity if a region is partitioned into sub-regions. It is argued in [19] we should partition a region if and only if the extra effort we need make is less than the complexity decrease. A *singleness* predicate $sngl$ is proposed to indicate whether a region R should be retained rather than partitioned

$$sngl(R) = \left[\forall\{R_i\}, D(R) - \sum D(R_i) \leq T(\{R_i\})\right] \tag{1}$$

where $\{R_i\}$ is a partition configuration of R, $D(\cdot)$ is the region description complexity, and T is the effort to trace boundaries between the regions in $\{R_i\}$. According to LEP, a human subject will finish his segmentation task by solving

$$\min_{\mathcal{S}} \ \lambda L\left(\mathcal{S}\right), \text{s.t.} \ \forall \ R \in \mathcal{S}, sngl(R) = \boldsymbol{TRUE}. \tag{2}$$

where $\mathcal{S}$ is the segmentation configuration in terms of a set of regions, and $L\left(\mathcal{S}\right)$ is the total tracing load of all boundaries induced by $\mathcal{S}$. The paper uses a bottom-up merging process to inversely simulate the top-down human segmentation process [19], and the corresponding merging criterion is called the *unit merging cost*[1]

$$umc(\{R_i\}) = \frac{D(\cup_i R_i) - \sum_i D(R_i)}{L(\{R_i\})}. \tag{3}$$

The paper solves a restricted version of eq. 2 under the *hierarchy* and *monotonicity* constraints, by repeatedly merging a pair of regions of the currently minimum *umc*. The corresponding segmentation algorithm exhibits the good boundary and region performance, together with the high efficiency.

However, it is not well-investigated in [19] how to accommodate the multilevel cues in the merging criteria, not to mention the emerging dense semantic predictions. The Fully Convolutional Neural Networks (FCNN) are proposed in [33], then widely used in semantic segmentation tasks. The FCNNs take the pixels as the input, and follow the basic structures of the standard convolutional neural networks which succeed in image classification challenges. Furthermore, the FCNNs can deal with images of arbitrary sizes, and gain flexible granularities using multiple processing streams [33]. The corresponding semantic segmentation methods can produce pixel-wise semantic labels, to say, maps of the probabilities of belonging to specific categories, see Fig. 1. They deal with the images without any preprocessing stages, and the results can be improved with the existing tools such as conditional random fields [34][36][37]. The FCNN-based methods obtain the state-of-the-art on several popular challenges such as the Pascal 2012 [41] and the MS COCO [45].

Compared with the poor capacity of their precursors [20][24], the FCNN-based approaches can deal with tens or even hundreds of categories. This meets the situation of image segmentation well, providing us the better support to handle the images in the wild. Nevertheless, it should be mentioned that the boundary accuracy is not an adherent

[1] The original term is *merging unit cost* in [19]. Here we change the order of the terms slightly for easier understanding.

objective for most semantic segmentation approaches, although the manually-annotated groundtruths are indeed composed of well-shaped regions. On the contrary, the region boundaries are emphasized in image segmentation tasks, thus it is interesting to observe whether we can benefit semantic segmentations on region boundaries in return.

3 Utilizing Semantic Cues in Image Segmentation

In order to find out the role of semantic cues in image segmentation, we turn to observing the human annotation processes. A good example is the widely-adopted image segmentation dataset, BSDS500. It is observed that human subjects are used to partition images hierarchically in the top-down style [1][43]: people partition images into large regions of distinct semantic categories, then divide them sequentially into sub-regions according to appearances, as shown in Fig. 3.

We use the bottom-up merging approach to simulate the human annotation processes inversely. In the final human segmentations, the pixels in each sub-region have two labels indeed: semantic category and appearance type. Consequently, it is only with both two labels the right way to discriminate pixels of different sub-regions. Therefore in the bottom-up machine processes, we need to use semantic cues at the very beginning, together with appearance cues to group pixels into appropriate sub-regions.

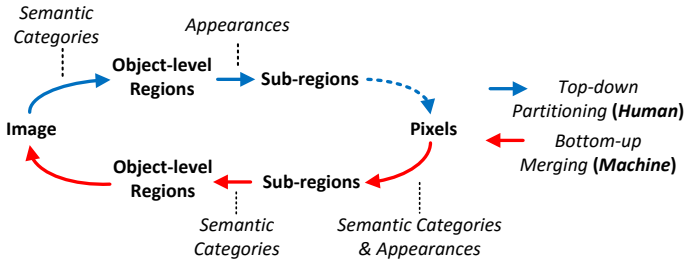

Fig. 3. The major supportive cues in different stages. Semantic cues are not essential for humans to partition category-level regions into sub-regions of different appearances, while they are essential for machines to discriminate pixels of different categories when merging them into sub-regions. The stage of the dotted arrow in blue seldom happens, but we present it here for the completeness considerations.

We notice that people are used to describe regions with the combinations of both appearances and semantic categories, such as *a dark dog* and *a curtain of black-white squares*, as shown in Fig. 4. It means the region description complexity is the sum of both appearance and semantic terms. We group colors and textures into clusters, then calculate discrete entropies as the color or texture description complexities for individual regions. A regularity term is designed for regions of slowly-varying appearances. It is the weighted sum of contour cues and regularized color distances between neighbor

pixels. The regularized color distances is inspired by the probabilistic density of gradients [46]. Then, we transform semantic cues into complexities and add them to the region description complexity $D(\cdot)$ in (1)

$$D(R) = w_c D_c(R) + w_t D_t(R) + w_s D_s(R) + w_r D_r(R) \tag{4}$$

where D_c, D_t, D_s, D_r are color, texture, semantic and regularity description complexities in turn, and w_c, w_t, w_s, w_r are weights which need to be trained.

Fig. 4. Image and pieces of single things. Left: the image and groundtruth boundaries. From middle left to the right: four pieces. The involved single things are described on the top of the pieces. Note: the image and groundtruth are from the BSDS500 train set.

If semantic cues are accurate enough, we can directly pick up pixels according to their semantic labels to assemble category-wise groups. Then it seems there is not space to gain improvements with image segmentations in hand. However in practice, due to the huge difficulty of semantic prediction tasks, there is much noise in cues from even FCNN-based approaches, especially those nearby the region boundaries, such as the feet and tail in the horse map in Fig. 1. Low level cues such as colors and textures are more effective in locating the boundaries [47]. Thus we hypothesize image segmentations are useful to correct mistakes in semantic cues to get accurate boundaries and consequently sensible regions. As a comparison, small superpixels can not provide equal support as they have lots of fake boundaries.

3.1 The Framework

Here's the framework utilizing dense semantic cues in Alg. 1:

Initially, each pixel is regarded as a region and put into a group according to the semantic map. Then we keep merging the region pairs of minimum umc's in each loop until the unit merging cost reaches the threshold λ_τ. Each time two region merge, the new merging corners are estimated and the corresponding tracing load is updated. We also use the postponed updating strategy in [19] to reduce wasted calculations in step 11 when updating umc's. All umc's and region pairs are stored in a direct access table [48], so we can search, insert or remove a record in the constant time. Along with the region merging, all interested category-wise groups are dynamically assembled. Finally we obtain the image segmentation $\mathcal{S}$ together with groups $\{G_l\}$ of specific categories.

For each pixel p, the semantic map $\mathcal{M}$ stores a vector $\langle v_{p,1}, v_{p,2}, ..., v_{p,N} \rangle$ measuring its memberships of semantic categories in $\mathcal{L} = \{l_1, l_2, \cdots, l_N\}$. We define the

Algorithm 1: Segmentation using Dense Semantic Cues

Input: Image I, semantic map $\mathcal{M}$ for labels in $\mathcal{L}$, threshold λ_τ
Output: Segmentation $\mathcal{S}$, category-wise groups $\{G_l\}$.

$S \leftarrow \{\langle p_i, l_i\rangle\}, G_l \leftarrow \phi$; // each pixel p_i as a region
foreach $p_i \in I$ **do**
  $l_i \leftarrow \arg\min_l D_s(p_i|l, \mathcal{M})$; // initial label of pixel p_i
  $G_{l_i} \leftarrow G_{l_i} \cup \{p_i\}$ // put pixels in initial groups for individual categories
end
$\langle R_a, R_b\rangle \leftarrow \arg\min umc(p_i, p_j), \lambda \leftarrow umc(R_a, R_b)$;
// the minimum *umc* on all pairs of adjacent pixels
while $\lambda < \lambda_\tau$ **do**
  $R_{ab} \leftarrow R_a \cup R_b, l_{ab} \leftarrow \arg\min_l D_s(R_{ab}|l, \mathcal{M})$;
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{R_{ab}\} - \{R_a, R_b\}, G_{l_a} \leftarrow G_{l_a} - R_a, G_{l_b} \leftarrow G_{l_b} - R_b,$
  $G_{l_{ab}} \leftarrow G_{l_{ab}} \cup R_{ab}$; // updating segmentation and category-wise groups
  collect new corners and update related tracing loads;
  update $umc(R, R_{ab})$ for each R adjacent to R_{ab} by (13);
  $\langle R_a, R_b\rangle \leftarrow \arg\min umc(R_i, R_j), \lambda \leftarrow umc(R_a, R_b)$;
  // the minimum *umc* on all pairs of adjacent regions in $\mathcal{S}$
end

semantic description complexity of a region R on category l as

$$D_s(R|l, \mathcal{M}) = -\sum_{p\in R} \log_2 v_{p,l}. \tag{5}$$

Each region has its most sensible semantic category. We let it be the category of the minimum description complexity

$$l_R = \arg\min_l D_s(R|l, \mathcal{M}), \tag{6}$$

and the semantic description complexity in (4) be

$$D_s(R|\mathcal{M}) = \min_l D_s(R|l, \mathcal{M}). \tag{7}$$

We omit $\mathcal{M}$ in $D_s(\cdot)$ for simplicity in the below.

4 Handling the Semantic Insufficiency

Most FCNN-based semantic models focus on hundreds of object categories at most, and all objects of other categories are treated as the *background*. According to eq. 6 and 7, d_s defined on two *background* regions is always zero, thus incapable of discriminating regions of unknown categories. We call it the semantic insufficiency problem.

For all pixels, their semantic description complexities should be much small on *correct* categories but large on *wrong* categories, as shown in Fig. 5. Thus we roughly let the small complexities equal to a small constant γ_1, and large complexities equal to

a large γ_2. For two *background* regions from different categories, suppose we have an enhanced semantic map $\mathcal{M}_e$ which is aware of their categories, then we have

$$D_s(R_1|\mathcal{M}_e) \sim \gamma_1|R_1|, D_s(R_2|\mathcal{M}_e) \sim \gamma_1|R_2|. \tag{8}$$

where $|\cdot|$ is the cardinality. Usually, the semantic category of the union of two regions comes from the previous categories. Hence

$$\begin{aligned} D_s(R_1 \cup R_2|\mathcal{M}_e) &\sim \min\left(\gamma_1|R_1| + \gamma_2|R_2|, \gamma_2|R_1| + \gamma_1|R_2|\right) \\ &= \gamma_1 \max(|R_1|, |R_2|) + \gamma_2 \min(|R_1|, |R_2|). \end{aligned} \tag{9}$$

Let $d_s = D_s(R_i \cup R_j) - D_s(R_i) - D_s(R_j)$, then we can estimate d_s with a new term

$$d_b \sim (\gamma_2 - \gamma_1) \cdot \min\left(|R_1|, |R_2|\right). \tag{10}$$

It means d_b is linearly related to the minimum region size, as shown in Fig. 5.c-d. However up to now, we do not really known whether the underlying categories are different. Suppose the *true* categories are different with a probability η, we use the expected value of d_b instead of random guesses

$$d_b = \begin{cases} \eta \cdot (\gamma_2 - \gamma_1) \min\left(|R_1|, |R_2|\right), & l_{R_1} = l_{R_2} = \textbf{\textit{BKG}}, \\ 0, & \text{otherwise}, \end{cases} \tag{11}$$

where ***BKG*** represents the *background* category. Inspired by the recent object proposal researches [17], we use location-sensitive prior probabilities to capture η in (11), please see the supplemental materials for details. It is possible to gain more improvements if we use complicated proposal models as in [49].

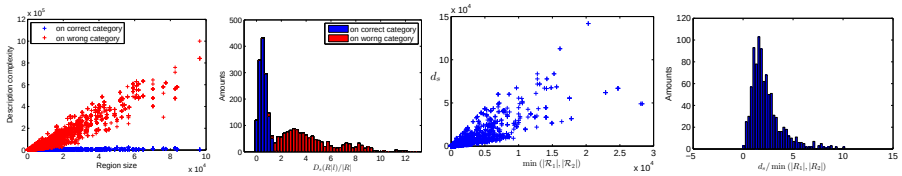

(a) (b) (c) (d)

Fig. 5. Semantic description complexities and region sizes. (a) $D_s(R|l)$ and region sizes. For each region R, we collect two samples by letting l be correct and wrong categories respectively. (b) Histogram of the ratio of $D_s(R|l)$ to region sizes. (c) d_s and the minimum region sizes. (d) Histogram of the ratio of d_s to the minimum region sizes. Most ratios are close to 2.2. Note: All samples are collected from the BSDS500 train set. Zoom in for better views.

5 Handling the Noise in Multilevel Cues

The noise in inaccurate semantic cues might tamper the accuracy of region boundaries as in Fig. 6.d-e. Besides this, the clustering accuracies of colors and textures have key

effect to umc calculations, but hard clustering methods such as K-Means might separate much similar colors/textures into different clusters. For example in Fig. 6.c, pixels in the sky region are grouped into five clusters, consequently it makes the following segmentation process divide the whole sky into several pieces. Soft clustering methods such as the Gaussian Mixture Model (GMM) use soft decision boundaries and should be capable of reducing the noise. However we find in practice its effect is far from expected. A recent work involves spatial continuity in color clustering by embedding them into manifolds, and obtains promising results [8]. Nevertheless it is too complicated to compute easily. Finally we choose to use hard clustering, and handle the potential multilevel noise jointly.

The cues such as inter-pixel color distances and local contours, are more accurate on locating boundaries. Intuitively, if two adjacent pixels have the same colors, they should originate from the same semantic object. So we use a logistics function with respect to the regularity term D_r as a soft switch σ, to determine whether the appearance/semantic cues should weigh in:

$$\sigma(R_i, R_j) = \left(1 + e^{-\alpha\left(\frac{d_r}{L(\{R_i,R_j\})} - \beta\right)}\right)^{-1} \tag{12}$$

where $d_r = D_r(R_i \cup R_j) - D_r(R_i) - D_r(R_j)$. All these parameters will be trained. Now the unit merging cost turns into

$$umc(R_i, R_j) = \frac{\sigma\left(w_c d_c + w_t d_t + w_s d_s\right) + w_u d_b + w_r d_r}{L(\{R_i, R_j\})}. \tag{13}$$

where d_c, d_t are defined in the same way as d_r and d_s. We do not handle the noise in d_b as it is an approximation. The experiments show the noise in multilevel cues can be suppressed effectively in this way, see Fig. 6.

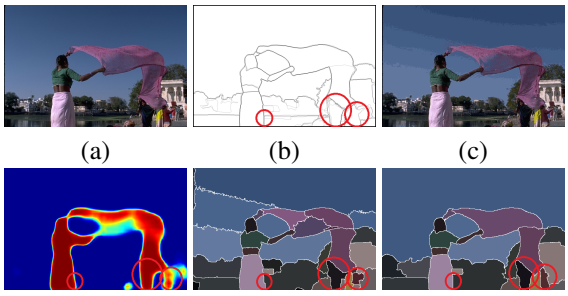

Fig. 6. Benefit of handling noise. (a) Image. (b) Averaged contour map of multiple groundtruths. (c) Clustered colors. Pixels are shown in mean colors. (d) Semantic map of the *person* category. Some noisy regions are bounded by red circles. Segmentations (e) without noise handling, and (f) with noise-handling.

The trick proposed in this section makes it possible to adopt simple but fast clustering methods to extract appearance cues, meanwhile without tampering the segmentation quality much. In fact, we adopt the Principal Components Analysis (PCA) based clustering in the experiments, please see the supplementary materials for details.

6 Experiments

We aim to verify whether semantic cues can benefit image segmentations as expected, and whether we can bring some improvement to semantic segmentations in return. There are two groups of experiments: image segmentation on the BSDS500 [11], and semantic segmentation on the Pascal VOC12 [41]. All parameters are trained on the BSDS500 train set, with the strategy mixing grid searching and gradient descending as in [16]. The semantic cues are obtained by the FCNN-based dense prediction method in [34]. There are only 20 meaningful categories involved as the model is established for the Pascal VOC. All source code and results would be open upon acceptance for easy reproducibility.

6.1 Image Segmentation

Table 1. Results on the BSDS500 test set. †: GPU time. We use structured edges [9] in *Ours*, and HED [13] (the original version with F=0.782) in *Ours-HED*. The lower VOI values are better.

	F_{op}		Covering		PRI		VOI		F			Time
	ODS	OIS	ODS	OIS	ODS	OIS	ODS	OIS	ODS	OIS	AP	
Human	0.56	0.56	0.72	0.72	0.88	0.88	1.17	1.17	0.80	0.80	-	-
EGB [4]	0.16	0.24	0.52	0.57	0.80	0.82	2.21	1.87	0.61	0.64	0.56	<1.0s
NCut [25]	0.21	0.27	0.45	0.53	0.78	0.80	2.23	1.89	0.64	0.68	0.45	600s+
MShift [5]	0.23	0.29	0.54	0.58	0.79	0.81	1.85	1.64	0.64	0.68	0.56	600s+
UCM [16]	0.35	0.38	0.59	0.65	0.83	0.86	1.69	1.48	0.73	0.76	0.73	240s
ISCRA [18]	0.35	0.42	0.59	0.66	0.82	0.85	1.60	1.42	0.72	0.75	0.46	240s+
MCG [17]	0.38	0.43	0.61	0.66	0.83	0.86	1.57	1.39	0.75	0.78	0.76	20s+
LEP [19]	0.42_{16}	0.46_{53}	0.62_{65}	0.68_{55}	0.83_{64}	0.86_{84}	1.46_{76}	1.29_{13}	0.75_{75}	0.79_{28}	0.81_{78}	1.0s
Ours	0.43_{09}	0.48_{61}	0.64_{30}	0.70_{12}	0.84_{26}	0.87_{76}	1.42_{00}	1.23_{19}	0.75_{94}	0.80_{03}	0.82_{69}	0.8s+0.4s†
Ours-HED	0.44_{37}	0.50_{14}	0.66_{16}	0.70_{73}	0.85_{52}	0.87_{85}	1.35_{61}	1.19_{71}	0.79_{38}	0.81_{60}	0.84_{99}	0.8s+0.7s†
SE [9]	-	-	-	-	-	-	-	-	0.75	0.77	0.80	0.4s
HED [13]	-	-	-	-	-	-	-	-	0.78	0.80	0.83	0.4s†

We adopt four region-based measures including F_{op} [50], Covering [16], PRI [16], VOI [16], and a boundary-based measure F [16]. Performance is considered better on larger measures except for VOI. We collect the performance on Optimal Dataset Scales (ODS) and Optimal Image Scales (OIS) as suggested in [16]. Our counterparts include seven methods, MCG [17], UCM [16], MShift [5], NCut [25], ISCRA[18], EGB [4], and LEP [19]. They are of the state-of-the-art performance or widely used nowadays.

For the sake of fairness, we use the same contour cues as adopted in MCG and LEP, structured edges [9]. All evaluation results are shown in Table 1 and Fig. 7. Our method

performs noticeably better on all region-based measures than its counterparts, indicating the semantic cues perform as expected to prevent merging regions of distinct categories. Especially in Fig. 7.b, on the large recall rates, our F_{op} values are hetergeneously better than that of LEP, indicating that the semantic cues weigh in even for small regions.

On the other hand, there are only margin improvements on the boundary-based measure. It implies that semantic cues have no significant effect on contour detections, as suggested in [47]. Another reason is the limited number of categories: only a small set of objects in BSDS500 have determined semantic categories. Meanwhile, some improvements are in small scales thus do not reflect much in measures. Nevertheless, they are remarkable in subjective evaluations, such as the dog mouth and the dancer leg in Fig. 8.

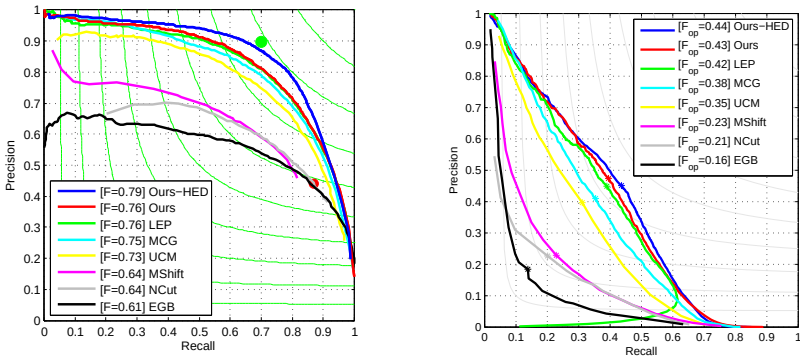

(a) (b)

Fig. 7. Performance on the BSDS500 test set. (a) F and (b) F_{op}.

Interestingly, in case that semantic cues are not accurate enough, we can also obtain high quality segmentations, as in the first group and last three groups in Fig. 8. It means that semantic cues and low-/mid-level cues can correct the mistakes in their companions. Furthermore, by utilizing the cutting-edge contour cues [13], we find the performance on all measures can be improved remarkably. It indicates that our framework is highly flexible to accommodate cues from various sources.

6.2 Semantic Segmentation

We use *ours-HED* in this part, then compare the segmentations with the original results, and those refined with fully connected CRFs [38] in [34]. We use the mean Intersection-over-Union (IoU) as the measure [41]. The threshold parameter λ_τ is different from that of the BSDS500 as the images have much larger sizes. We test all methods on the Pascal VOC12 val set, and the results are shown in Table 2.

The performance of our method is noticeably better than that of [34]. It proves our framework do bring some improvements to the original semantic segmentations. Especially, our advantages are more remarkable on objects of big connected regions

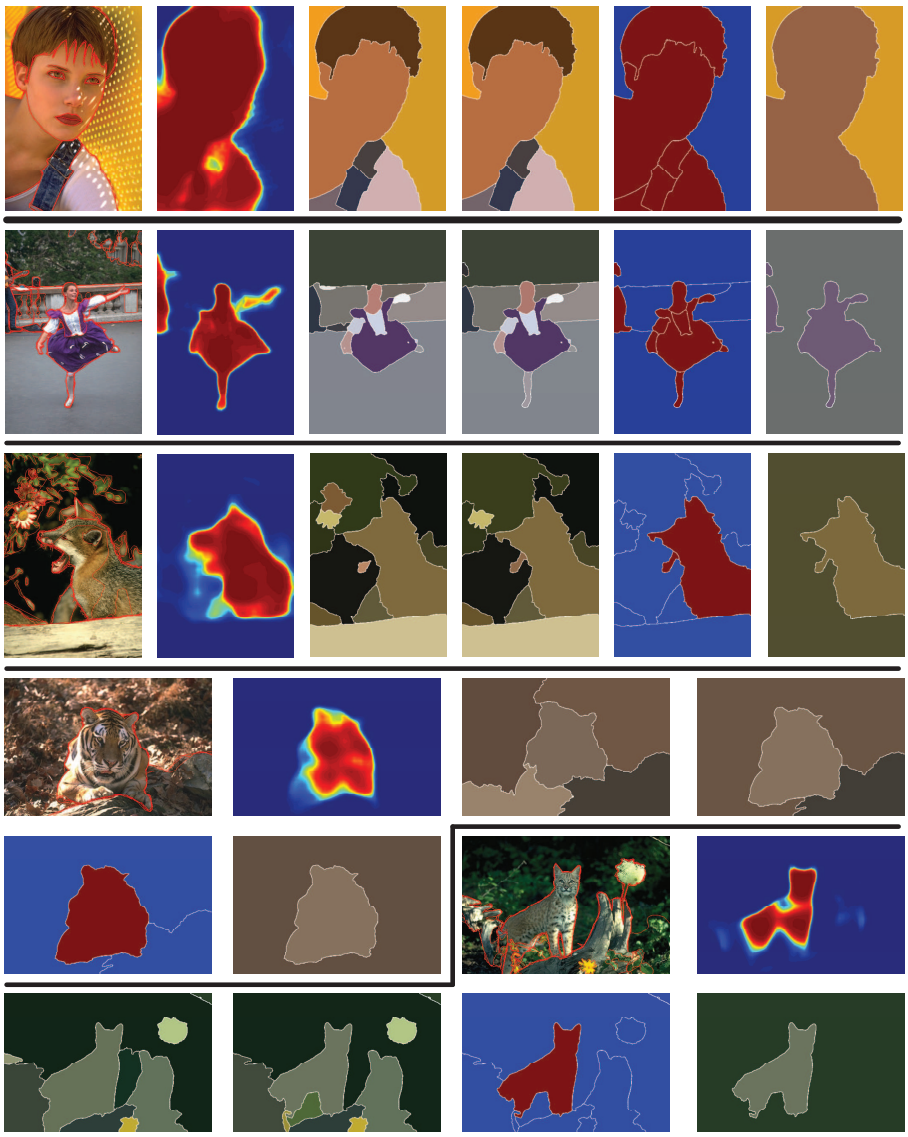

Fig. 8. BSDS500 images and segmentations. In each group: image (covered by groundtruth boundaries in red), semantic map of the major category, LEP segmentation, Ours-HED segmentation, category-wise subregion map, and semantic segmentation. Note: all results are ODS ones. Zoom in for better views.

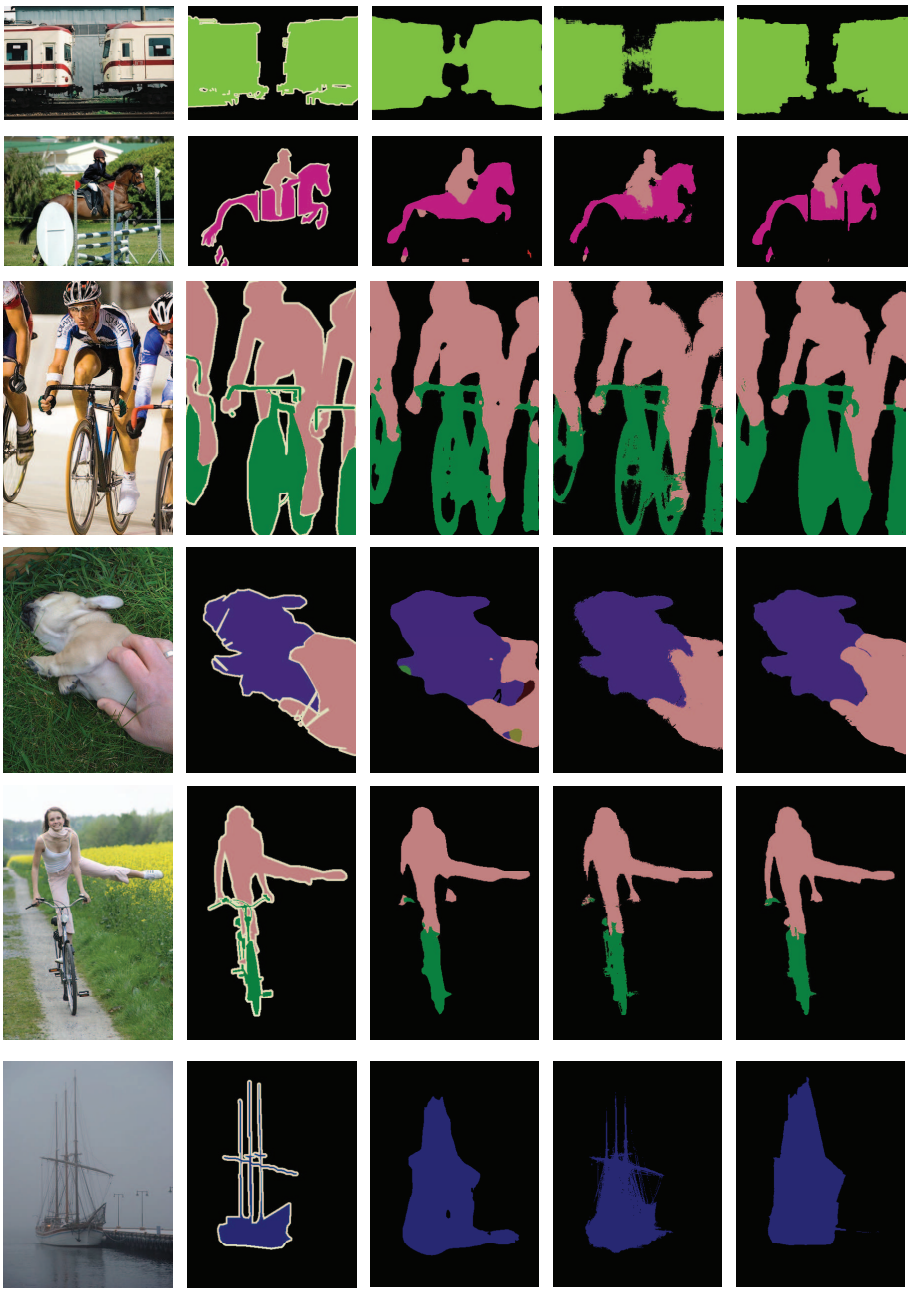

Fig. 9. Pascal VOC images and segmentations. Left: image. Middle left: groundtruth. From middle to right: segmentations of Deeplab, Deeplab+CRF, and Deeplab+*Ours-HED*. Note: all results are ODS ones. Zoom in for better views.

Table 2. IoUs (%) on the Pascal VOC12 val set. The best values are shown in bold.

	plane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
Deeplab [34]	86.7	58.9	85.8	76.3	78.4	91.7	85.7	87.6	41.3	88.3	63.4	84.9	82.0	81.6	84.3	65.4	85.6	63.6	88.0	68.7	78.2
Deeplab+CRF [34]	91.1	**61.0**	90.9	80.9	82.2	92.9	87.0	91.4	45.9	92.3	**67.1**	89.1	85.8	83.8	86.7	**71.7**	91.0	68.8	**89.7**	73.9	81.8
Deeplab+Ours-HED	**91.4**	55.8	**92.3**	**81.3**	**85.0**	**93.1**	**87.3**	**92.2**	**46.9**	**93.3**	**67.1**	**91.9**	**87.0**	**83.9**	**87.5**	70.1	**92.4**	**70.0**	89.0	**75.3**	**82.3**

such as bottles and dogs, however less on those of long thin structures such as the ship in Fig. 9. The reason is, our method is more inclined to generate connected regions without holes, while CRF-based methods are much good at delicate structures. The style of our results looks more like that of human annotations, given that there are less noise-like structures than in CRF results. The boundaries become smoother and more accurate as expected. It makes our results more remarkable in subjective evaluations than in IoU measures, see Fig. 9.

7 Discussion and Conclusions

The experiments indicate that our method obtains the better segmentation performance, and the semantic cues work as expected from the very beginning. It implies that the way we exploit semantic cues in image segmentation is effective and sensible. Besides this, the framework is flexible so that we can adopt the cues from various sources. The way we propose to handle the noise in multilevel cues is effective in improving the robustness, making it possible to adopt simple color/texture clustering methods for the higher efficiency. We can replace them with more delicate ones for the better performance. At last, we believe a more comprehensive semantic model involving plenty of categories would benefit the segmentations substantially.

Meanwhile it is easy to extract semantic segmentation results. Compared with the original results, our method gains the better performance together with more accurate region boundaries. Thus it is a good choice to refine the dense semantic segmentations.

References

1. Martin, D., Fowlkes, C., Tal, D., Malik., J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statics. In: ICCV. (2001) 1, 3, 5
2. Cremers, D., Rousson, M., Deriche, R.: A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. In: IJCV. (2007) 1
3. Bagon, S., Boiman, O., Irani, M.: What is a good image segment? a unified approach to segment extraction. In: ECCV. (2008) 1, 4
4. Felzenszwalb, P., Huttenlocher, D.: Efficient graph-based image segmentation. IJCV (2004) 1, 10
5. Comaniciu, D., Meer, P.: Mean shift: a robust approach toward feature space analysis. IEEE TPAMI (2002) 1, 10
6. Ma, Y., Derksen, H., Hong, W., Wright, J.: Segmentation of multivariate mixed data via lossy data coding and compression. IEEE TPAMI (2007) 1

7. Yarkony, J., Ihler, A., Fowlkes., C.: Fast planar correlation clustering for image segmentation. In: ECCV. (2012) 1
8. Yu, Y., Fang, C., Liao, Z.: Piecewise flat embedding for image segmentation. In: ICCV. (2015) 1, 3, 9
9. Dollar, P., Zitnick, C.: Fast edge detection using structured forests. IEEE TPAMI (2015) 1, 10, 11
10. Donoser, M., Schmalstieg, D.: Discrete-continuous gradient orientation estimation for faster image segmentation. In: CVPR. (2014) 1
11. Martin, D., Fowlkes, C., Malik, J.: Learning to detect natural image boundaries using local brightness, color, and texture cues. IEEE TPAMI (2004) 1, 10
12. Ren, X., Malik, J.: Learning a classification model for segmentation. In: ICCV. (2003) 1
13. Xie, S., Tu, Z.: Hlistically-nested edge detection. In: ICCV. (2015) 1, 10, 11
14. Kokkinos., I.: Pushing the boundaries of boundary detection using deep learning. In: ICLR. (2016) 1
15. Arbeláez, P.: Boundary extraction in natural images using ultrametric contour maps. In: CVPR Workshop. (2006) 1
16. Arbeláez, P., Maire, M., Fowlkes, C., Malik., J.: Contour detection and hierarchical image segmentation. IEEE TPAMI (2011) 1, 2, 3, 10
17. Arbeláez, P., Pont-Tuset, J., Barron, J., , Marques, F., Malik, J.: Multiscale combinatorial grouping. In: CVPR. (2014) 1, 3, 8, 10
18. Ren, Z., Shakhnarovich, G.: Image segmentation by cascaded region agglomeration. In: CVPR. (2013) 1, 10
19. Zhao., Q.: Segmenting natural images with the least effort as humans. In: BMVC. (2015) 1, 2, 3, 4, 6, 10
20. Maire, M., Yu, S., Perona., P.: Object detection and segmentation from joint embedding of parts and pixels. In: ICCV. (2011) 1, 2, 4
21. Russell, B.C., Freeman, W.T., Efros, A.A., Sivic, J., Zisserman., A.: Using multiple segmentations to discover objects and their extent in image collections. In: CVPR. (2006) 1
22. Tu, Z., Chen, X., Yuille, A.L., Zhu, S.C.: Image parsing: Unifying segmentation, detection, and recognition. (2005) 1
23. Wu, J., Zhu, J.Y., Tu, Z.: Reverse image segmentation: a high-level solution to a low-level task. In: BMVC. (2014) 1, 2
24. Borenstein, E., Ullman., S.: Combined top-down/bottom-up segmentation. (2008) 2, 4
25. Shi, J., Malik, J.: Normalized cuts and image segmentation. IEEE TPAMI (2000) 2, 10
26. Brox, T., Bourdev, L., Maji, S., Malik., J.: Object segmentation by alignment of poselet activations to image contours. In: CVPR. (2011) 2
27. Parkhi, O., Vedaldi, A., Jawahar, C.V., Zisserman., A.: The truth about cats and dogs. In: ICCV. (2011) 2
28. Yang, Y., Hallman, S., Ramanan, D., Fowlkes., C.: Layered object detection for multi-class segmentation. In: CVPR. (2010) 2
29. Carreira, J., Sminchisescu., C.: Constrained parametric min-cuts for automatic object segmentation. In: CVPR. (2010) 2
30. Uijlings, J., van de Sande, K., Gevers, T., Smeulders., A.: Selective search for object recognition. In: IJCV. (2013) 2
31. Arandjelović, R., Zisserman., A.: Smooth object retrieval using a bag of boundaries. In: ICCV. (2011) 2
32. Ion, A., Carreira, J., Sminchisescu., C.: Probabilistic joint image segmentation and labeling. In: NIPS. (2011) 2
33. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. (2015) 2, 4

34. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille., A.: Semantic image segmentation with deep convolutional nets and fully connected crfs. In: ICLR. (2015) 2, 4, 10, 11, 14
35. Dai, J., He, K., Sun, J.: Boxsup: exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: ICCV. (2015) 2
36. Lin, G., Shen, C., Reid., I.: Efficient piecewise training of deep structured models for semantic segmentation. In: arXiv. (2015) 2, 4
37. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr., P.: Conditional random fields as recurrent neural networks. In: arXiv. (2015) 2, 4
38. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. In: NIPS. (2011) 2, 11
39. Liu., Z., Li, X., Luo, P., Tang, X.: Semantic image segmentation via deep parsing networks. In: ICCV. (2015) 2
40. Fu, X., Wang, C., Chen, C., Wang, C., Kuo., C.: Robust image segmentation using contour-guided color palettes. In: ICCV. (2015) 3
41. Everingham, M., Eslami, S.M.A., van Gool, L., Williams, C., Winn, J., Zisserman., A.: The pascal visual object classes challenge: a retrospective. (2015) 3, 4, 10, 11
42. Zipf, G.: Human behaviour and the principle of least effort. Addison-Wesley Press (1949) 3
43. BerkeleyVisionGroup: The berkeley segmentation dataset and benchmark. (http://www.eecs.berkeley.edu/Research/ Projects/CS/vision/grouping/segbench/) 3, 5
44. Rao, S., Mobahi, H., Yang, A., Sastry, S., Ma., Y.: Natural image segmentation with adaptive texture and boundary encoding. In: ACCV. (2009) 4
45. Lin, T., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C., Dollar., P.: Microsoft coco: Common objects in context. In: ECCV. (2014) 4
46. Jia, Y., Darrell., T.: Heavy-tailed distances for gradient based image descriptors. In: NIPS. (2011) 6
47. Zitnick, C.L., Parikh, D.: The role of image understanding in contour detection. In: CVPR. (2012) 6, 11
48. Cormen., T., Leiserson, C., Rivest, R., Stein., C.: Introduction to algorithms. MIT Press (2009) 6
49. Wang, C., Zhao, L., Liang, S., Zhang, L., Jia, J., Wei., Y.: Object proposal by multi-branch hierarchical segmentation. In: CVPR. (2015) 8
50. Pont-Tuset, J., Marques, F.: Measures and meta-measures for the supervised evaluation of image segmentation. In: CVPR. (2013) 10